

Delivering the AI Strategy - the use of new AI technologies in industry and the public sector

Date and Location: 23rd February 2022 at The Royal Society

Chair:	John Neilson Trustee, The Foundation for Science and Technology
Speakers:	Professor Dame Wendy Hall DBE FRS FREng Regius Professor of Computer Science, University of Southampton Lord Clement-Jones CBE House of Lords Professor Geraint Rees FMedSci Pro-Vice-Provost, AI, University College London Professor Tom Rodden Chief Scientific Adviser, Department for Digital, Culture, Media and Sport
Report Author:	Alex Elliott Foundation Future Leader, Cranfield University
Sponsor:	UKRI Trustworthy Autonomous Systems Hub
Audio/Video Files:	www.foundation.org.uk
Hash tag:	#fstAI . Twitter Handle: @FoundSciTech .

JOHN NEILSON, commenced proceedings by welcoming participants both in-person and online, thanking the event sponsors – the UKRI Trustworthy Autonomous Systems Hub – and providing context for the national AI Strategy¹, published by the government in September 2021.

PROFESSOR DAME WENDY HALL gave the first talk. She initially focused on her role as co-chair of the government’s AI Review², published in October 2017, and subsequent positions on the AI Council and as AI Skills Champion.

Dame Wendy first presented a summary of the journey from the initial AI Review to the current AI Strategy and Office for AI, who were presented the AI Roadmap³ by the AI Council in January 2021. This roadmap

formed the basis of the AI Strategy, which is motivated by its vision for the UK “to remain an AI and science superpower fit for the next decade.” Particular emphasis was placed on the inclusion of the word “remain”, given the UK’s current position of third in the world for AI, vastly outperforming its population size, as a result of our legacy of “doing AI since before it was called AI”. The outcomes of the Strategy will be across three pillars – “Investing in the long-term needs of the AI ecosystem”, “Ensuring AI benefits all regions and sectors”, “Governing AI effectively” – and all three permeated through the talks.

Exploring the AI Strategy overview further, Dame Wendy stressed the importance of sustained investment in AI, if the UK is to maintain its status as an AI superpower and to achieve the goals of these three pillars. The ability to build upon existing strengths in academic research and fostering novel start-ups and allowing them to scale up in the UK represents a key challenge for the coming decade. It is imperative that solutions to this challenge are created in a manner that

1 <https://www.gov.uk/government/publications/national-ai-strategy>

2 <https://www.gov.uk/government/publications/growing-the-artificial-intelligence-industry-in-the-uk>

3 <https://www.gov.uk/government/publications/ai-roadmap>

is good for people, for society, for business, and for government, as well as constituting good public value for money and contributing to the growth of GDP.

More specifically, key elements of the strategy include data trusts (with achieving an ethos of public trust in AI recognised as an important goal), driving adoption of AI across sectors (e.g. through the NHS AI Health and Care Award), publishing guidelines (e.g. for ethics, bias in algorithmic decision making, machine learning with limited data), and strategic cooperation with international partners (such as through the GPAI).

To conclude, Dame Wendy summarised the current successes and future opportunities from the R&D funding component of the AI ecosystem, with the UKRI and Alan Turing Institute – the National Institute for Data Science and AI – identified as the pivotal figures in this process. The £46 million Turing AI Fellowship programme was used to attract and retain exceptional researchers, with plans to double the number of these as part of the AI Strategy. Finally, Dame Wendy announced the funding secured for 2000 elite AI Scholarships, with half of these specifically awarded to underrepresented groups – namely women, Black, and disabled students – and she closed with the pertinent remark that “if it isn’t diverse, it isn’t ethical.”

LORD CLEMENT-JONES started by highlighting the symmetries between his and Dame Wendy’s journey towards the AI Strategy, with the House of Lords AI Select Committee’s first report on AI in the UK published at a similar time to the AI Review. He went on to underline the definitional challenges associated with AI and to note that data quality is key to AI applications.

Lord Clement-Jones provided an overview of the AI policy ecosystem, including those bodies established as a direct result of AI Council recommendation. Cooperation and collaboration between the various members of this ecosystem will be vital to its success, as evidenced, for example, by the new Digital Regulation Cooperation Forum.

The focus then turned to the importance of ethical regulation for AI, which is to be informed by the AI Governance White Paper due in the coming months. Lord Clement-Jones was keen to underscore the fact that regulation is not necessarily the enemy of innovation and, in fact, it can act as a stimulus, particularly if it can be leveraged to engender public trust in AI. Quoting the title of Dr Stephanie Hare’s new book, “technology is not neutral”, and it is

paramount that the purposes and implications of AI must be made clear. Referring to statements made by Professor Margaret Boden, as well as the ethical decisions made by the scientists in Jurassic Park, Lord-Clement Jones asked whether “even if AI could do something, should it?” Acknowledging this need for robust governance, the CDI has developed a roadmap for delivering international standards and an ecosystem of AI assurance tools, published guidelines on the application of AI in the public sector and on procurement, and established a pilot AI Standards Hub.

In contrast, Lord Clement-Jones identified potential decisions through which progress may be inhibited, including the post-Brexit decision to remove Article 22 of the UK GDPR (instead he called for an expansion to include automated decision making and inferred data), the removal of the requisite for Data Protection Officers and data protection impact assessments, and the withdrawal of support for the recommendations made by the Council of Europe’s ad hoc committee on AI. Moreover, more robust mechanisms for ethical and transparent automated decisions in the public sector – such as live facial recognition – are required, and should be encapsulated in a risk-based form of horizontal regulation.

PROFESSOR GERAINT REES opened his talk with an assertion that the success of the AI Strategy must focus on people, rather than technology, if it is to deliver strategic advantage. The Covid-19 pandemic provided valuable insight into our relationship with AI and its associated technologies as an important part of our shared future, though not one that can replace the inherent human need for social interaction. As recommended by the OECD, AI “must put people and planet first” and, by acknowledging the need for open science and the potential for misuse, it can be considered a force for good.

The future outlined in the AI Strategy will require fundamental changes to the workforce, as well as the skillset they possess. Although there will be a need for sufficient software engineering capabilities, Professor Rees placed greater importance on the need to master successful interactions between sectors, recognising the vital role of the arts and humanities in understanding the human experience. The challenge of defining effective partnerships to combine sectoral knowledge bases remains non-trivial, with the approach of bringing together individuals with complementary skills – utilising the experience from

academia in facilitating interdisciplinary discussions – providing greater promise than training individuals with skills in several fields.

Professor Rees's penultimate observation returned to the topic of fairness and bias, a topic that has only recently received recognition as a significant problem for the discipline of AI. Citing several key examples of machine learning applications that erroneously captured implicit biases from historical data, he went on to introduce the notion that some biases are important to distinguish, such as diseases with strong biological, environmental, and economic indicators. These complex medical diagnoses are representative of a wider set of problems in which the categorisation of data is far from being straightforward. Such challenges are not superficial, rather they are conceptual and fundamental, and cannot be overcome by ethical guidelines alone.

Finally, the grand challenge of achieving benefits from AI across the UK must be undertaken collaboratively, as outlined in the government roadmap on the subject. Professor Rees advocated for an AI ecosystem incorporating infrastructure, expertise, and entrepreneurship, and providing “the mixture of the right conditions” to allow both local and national prosperity.

PROFESSOR TOM RODDEN provided the final talk, addressing the drivers required to deliver the AI strategy. The UK's aforementioned position as a world leader in AI is an incredible asset, but a doubling down on the £2.3 billion invested thus far is required to ensure that this is maintained across the next decade. Professor Rodden reiterated the call made by Dame Wendy for the UK not to “take the foot off the gas” and remain adaptable as new challenges arise, highlighting the incidence of these messages in both the AI Council Roadmap and the AI strategy itself.

As UK academics and research bases continue to provide innovation, the transition of these solutions into real-world applications, capable of providing value across the country, must be accelerated. As AI moves from being a separate field towards integration across society, future developments will need to be a partnership between people and technological advances.

At the centre of the three pillars discussed above is the overarching theme of “Diversity and Public Trust”. The AI ecosystem, as discussed by all four speakers, will not simply emerge without stimulation from long-term investment, and Professor Rodden outlined the needs

that such a system will have to guarantee long term success. Chief among these was diversity at all levels, and he asked the key question “whose AI do we want?” Further needs were identified as the alignment of AI R&D funding (particularly with the wider ecosystem), the coordination of the wider computational needs (e.g. large-scale compute), and the link between open and machine-readable government datasets and AI models.

Professor Rodden went on to identify important steps to establish nationwide benefits from AI. These included targeted efforts to promote adoption in lower-AI-maturity sectors, increased capabilities for guaranteeing trustworthiness and transparency for AI technologies, and working with government agencies to determine applications in which AI can provide catalytic contributions to strategic challenges. At its core, collaborative working will be a key driver in ensuring widespread adoption, particularly for businesses without the capacity to establish internal AI departments.

After reiterating the need for effective AI governance, Professor Rodden provided a critical outline for the short-term steps required for implementing the AI Strategy. The Office for AI must continue its engagement with partners across government, industry, and academia to develop a coordinated approach to broaden and refine mechanisms for AI adoption. This process must carefully consider its interrelationship with international partners, with particular attention to be given to those multinational corporations whose reach goes beyond that of nation states.

DISCUSSION

The first question from the floor considered the successes, and potential shortcomings, of the current AI ecosystem. The panel began by highlighting that questions of this nature serve to illustrate the embedded nature of AI, noting John McCarthy's claim that, once it works, it is no longer referred to as AI. A host of specific AI successes were listed, including those within FinTech, EdTech, and agriculture, with the final statement on the matter advocating for the huge impact AI will have in meeting Sustainable Development Goals.

The discussion then turned to challenges of regulation and standardisation for such a rapidly-evolving area, with the latter again disadvantaged by the variation in meaning of the term “AI”. Following the rapid emergence of AI centres across government,

academia, and the wider world, there is a critical need for appropriate regulation and identification of adequate governance mechanisms. It is likely the case that a single regulator will not suffice and that domain-specific bodies – particularly in the case of medical data, for which informed consent cannot be given at the time of sample collection – may be required if the consent relationship is to be based on trust, rather than transaction. On the topic of effective governance and multinational businesses, the GPAI was identified as a useful form of collaboration, with previous there being numerous examples of such companies welcoming external regulation over that which is driven by industry itself.

Questions from online participants initially focused on skills, citing the claim, made in the Capital Economics report on AI Activity in the UK, that there may be a skills shortage of 50-100,000 in the future. The panel highlighted the need for experts remaining within their sectors and interacting with AI specialists, rather than simply transitioning large numbers of people into the field. Whereas there has been previous concern over the loss of jobs as a result of widespread adoption of AI, the new value propositions and supply chains this adds will instead result in a host of new positions and skills, starting at the apprentice level.

The panel finally discussed the intersection of AI and several of the major challenges currently facing the world. With regard to defence, international treaties, similar to those for nuclear proliferation, were advocated as a strategy for collaboration with international partners that may hold themselves to disparate ethical standards. Similarly, the climate challenge represents another area in which AI will prove beneficial, albeit with careful consideration required to overcome the compute and energy consumption associated with technologies such as Web 3.0 and blockchain. This message is one that is indicative of the entire evening's discussions: the field of AI represents an opportunity to massively improve in the UK, and beyond, and the AI Strategy will be key to ensuring that this is done safely, responsibly, and ethically.

Alex Elliott